

Three-dimensional wideband electromagnetic modeling on massively parallel computers

David L. Alumbaugh, Gregory A. Newman, Lydie Prevost, and John N. Shadid

Sandia National Laboratories, Albuquerque, New Mexico

Abstract. A method is presented for modeling the wideband, frequency domain electromagnetic (EM) response of a three-dimensional (3-D) earth to dipole sources operating at frequencies where EM diffusion dominates the response (less than 100 kHz) up into the range where propagation dominates (greater than 10 MHz). The scheme employs the modified form of the vector Helmholtz equation for the scattered electric fields to model variations in electrical conductivity, dielectric permittivity and magnetic permeability. The use of the modified form of the Helmholtz equation allows for perfectly matched layer (PML) absorbing boundary conditions to be employed through the use of complex grid stretching. Applying the finite difference operator to the modified Helmholtz equation produces a linear system of equations for which the matrix is sparse and complex symmetrical. The solution is obtained using either the biconjugate gradient (BICG) or quasi-minimum residual (QMR) methods with preconditioning; in general we employ the QMR method with Jacobi scaling preconditioning due to stability. In order to simulate larger, more realistic models than has been previously possible, the scheme has been modified to run on massively parallel (MP) computer architectures. Execution on the 1840-processor Intel Paragon has indicated a maximum model size of $280 \times 260 \times 200$ cells with a maximum flop rate of 14.7 Gflops. Three different geologic models are simulated to demonstrate the use of the code for frequencies ranging from 100 Hz to 30 MHz and for different source types and polarizations. The simulations show that the scheme is correctly able to model the air-earth interface and the jump in the electric and magnetic fields normal to discontinuities. For frequencies greater than 10 MHz, complex grid stretching must be employed to incorporate absorbing boundaries while below this normal (real) grid stretching can be employed.

Introduction

Great strides have been made over the last decade in geophysical electromagnetic (EM) modeling of three-dimensional (3-D) structures. These advances have been primarily due to the rapid improvements in computer speed and memory. For compact bodies residing in a layered Earth, integral equation (IE) solutions [Tripp and Hohmann, 1984; Newman *et al.*, 1986; Xiong, 1992] offer the most efficient way of producing a solution. However, as the complexity and size of the model grows, IE methods become numerically cumbersome.

Recently, a series of papers have appeared that employ various approximations to the integral

equations governing the EM fields [Habashy *et al.*, 1993; Torres-Verdin and Habashy, 1994; Zhdanov and Fang, 1995]. For general models these methods are much quicker than the full IE solution because they do not require the inversion of a large, dense matrix. However, the approximations often exhibit limitations in terms of the frequency that can be employed and the conductivity contrasts that can be simulated within the model. In addition, to this point these types of approximations have not been implemented in the radar range (greater than 10 MHz) where wave propagation becomes important. Thus to ensure accurate calculation of the EM fields for general models over a wide range of frequencies and material contrasts, a differential equation solution needs to be employed.

Two different forms of differential equation solutions exist: finite difference (FD) methods and finite element (FE) methods. The advantage of these

techniques over integral equations is that a linear system of equations results in which the matrix is very sparse. This property allows for the implementation of iterative Krylov subspace methods to solve the system which are much more efficient than direct matrix inversion. Although finite element techniques are more flexible in terms of the geometry of the mesh that can be employed, we have chosen to employ an FD scheme. With an FD scheme a well-structured grid results, which is acceptable for the models we plan to consider.

Even with these computationally efficient solutions, the complexity, and thus the realism of the models that can be simulated on traditional serial computers, is limited by memory and flop rate of the processor. However, with the development of massively parallel machines, the rate at which the simulations can proceed has dramatically increased because thousands of processors can operate on the problem simultaneously. In this paper we examine the implementation of a FD scheme on a parallel platform and demonstrate its usefulness over a wide frequency range for different types of geophysical scenarios. In addition we will briefly examine two different types of Krylov solvers and illustrate how the choice of the proper solver can have a dramatic impact on the solution time.

Theoretical Development

In order to simulate the EM response of a 3-D earth, we numerically solve the frequency domain version of the vector Helmholtz equation for the scattered electric fields using a finite difference approximation on a staggered grid (defined below) as outlined by *Yee* [1966]. The use of this type of grid has recently experienced a gain in popularity in EM geophysical applications; *Druskin and Knizhnerman* [1988, 1994], *Smith* [1992], *Wang and Hohmann* [1993], *Newman* [1995], and *Newman and Alumbaugh* [1995] all employ some type of staggered finite difference grid to solve for the EM fields in both the time and/or frequency domain.

The FD solution we shall outline has been designed to compute the 3-D EM response for a wide variety of Earth properties for frequencies ranging from approximately 100 Hz up to 100 MHz. This scheme is an extension of the one outlined by *Newman and Alumbaugh* [1995] with the major theoretical differences being (1) the incorporation of absorbing boundary conditions (ABCs) and (2) the

ability to model variations in magnetic permeability. The ABCs are required to simulate the response of frequencies greater than 10 MHz where wave propagation becomes dominant over lower-frequency EM diffusion as without them, erroneous results are produced. We have chosen to employ the “perfectly matched layer” (PML) absorbing boundary conditions originally developed by *Berenger* [1993] for two-dimensional (2-D) time domain calculations and later modified for 3-D by *Katz et al.* [1994] and *Chew and Weedon* [1994]. This method uses a modified form of the vector Helmholtz equation in which the absorption is incorporated through the use of complex grid stretching. Here we discuss the theoretical development of the finite difference solution and leave the discussion of the reasons for using the PML and its properties for a later section and Appendix A.

Assuming a time harmonic dependence of $e^{i\omega t}$ where $i = \sqrt{-1}$, incorporating variable magnetic permeability into equation (8) of *Newman and Alumbaugh* [1995] and using the same form for the modified differential operators as given by *Chew and Weedon* [1994] results in the following expression for the vector Helmholtz equation:

$$\nabla_h \times \frac{\mu_p}{\mu} \nabla_e \times \mathbf{E}^s = -i\omega\mu_p(\sigma + i\omega\epsilon)\mathbf{E}^s - i\omega\mu_p[(\sigma - \sigma_p) + i\omega(\epsilon - \epsilon_p)]\mathbf{E}^p - i\omega\mu_p \nabla \times \left[\left(\frac{\mu - \mu_p}{\mu} \right) \mathbf{H}^p \right] \quad (1)$$

where

$$\nabla_e = \hat{\mathbf{i}} \frac{1}{e_x} \frac{\partial}{\partial x} + \hat{\mathbf{j}} \frac{1}{e_y} \frac{\partial}{\partial y} + \hat{\mathbf{k}} \frac{1}{e_z} \frac{\partial}{\partial z} \quad (2)$$

and

$$\nabla_h = \hat{\mathbf{i}} \frac{1}{h_x} \frac{\partial}{\partial x} + \hat{\mathbf{j}} \frac{1}{h_y} \frac{\partial}{\partial y} + \hat{\mathbf{k}} \frac{1}{h_z} \frac{\partial}{\partial z}. \quad (3)$$

In these expressions the electrical conductivity, magnetic permeability, and dielectric permeability are denoted by σ , μ , and ϵ , respectively, where the “p” subscript designates a background or primary value; \mathbf{E}^s , \mathbf{E}^p , and \mathbf{E}^t are the scattered, primary, and total electric fields where $\mathbf{E}^t = \mathbf{E}^p + \mathbf{E}^s$; \mathbf{H}^s , \mathbf{H}^p , and \mathbf{H}^t are the associated magnetic fields; e_j and h_j for $j = x, y, z$ are coordinate-stretching variables which stretch the x , y , and z coordinates. As shown in Appendix A, when e_j and h_j are




Figure 1. The staggered grid for the Helmholtz equation for the electric fields. The dark circle at the center represents node (i, j, k) , which has the three designated components of the electric field assigned to it. The large arrows represent the 13 unknown electric field values needed to form the equation for E_x at node (i, j, k) , with the other arrows representing the additional fields needed to form the equations for E_y and E_z . The gray circles represent nodal points to which unknown electric field values are assigned that are needed to complete the three equations at node (i, j, k) ; the open circles represent nodal points from which no information is needed for these equations.

complex, then the medium is perfectly absorbing and the degree of absorption is independent of the angle of incidence. In addition, note that $(\mu - \mu_p)\mathbf{H}^p$ and $[(\sigma - \sigma_p) + i\omega(\epsilon - \epsilon_p)]\mathbf{E}^p$ can be thought of as equivalent source vectors for the background medium. The primary fields and background values found in these source vectors can be that of a whole space, a layered half space, or some previously run model for which the results were saved to be incorporated as the primary field. For all of the examples presented here, a whole space background has been assumed. As demonstrated in Appendix B, when (1) is discretized, a linear system of equations results in which the matrix is complex symmetrical.

The benefit of the above formulation over a total field formulation is that if the source is removed from zones of anomalous electromagnetic properties, then this equivalent source exhibits a smoother spatial dependence than that of an impressed dipole. Thus the fields are better behaved in the vicinity of the transmitter. It must be mentioned, however, that if the source is within a region of

anomalous properties, then problems can arise due to the rapid variation of the primary field within the equivalent source terms. This phenomenon will be addressed below.

Numerical Solution

The Finite Difference Approximation

To cast expression (1) into a system of linear equations, we discretize the earth (and the air above) into a mesh of rectangular blocks. Again, because we are employing the scattered field formulation, this grid does not need to be finely discretized near the source unless it is located within or near a region of anomalous electromagnetic parameters. To solve for \mathbf{E}^s as given in equation (1), we employ a staggered grid [Yee, 1966] for each component of the electric field, as shown in Figure 1. Following Wang and Hohmann [1993], this corresponds to sampling the electric fields on the edges of the cells and the magnetic fields at the center of the cell faces.

The above formulation requires that the admittiv-




Figure 2. Illustration to demonstrate how the average admittivity is calculated half way along a cell edge in the y direction, and the average magnetic permeability is calculated at the center of the cell face.

ity, defined as $\hat{y} = \sigma + i\omega\epsilon$, be computed halfway along a given cell edge in Figure 1. *Wang and Hohmann* [1993] showed that for low-frequency calculations, an average conductivity can be evaluated by tracing out a line integral of the magnetic field centered on the midpoint of the cell edge. Using their formulation, the resulting admittivity is simply the weighted sum of the admittivities of the four adjoining cells. For example, in Figure 2 the averaged admittivity in the y direction at $(i + 1, j + 1/2, k + 1)$ is given by

where $A_{i,j+(1/2),k}$ is a weighting function that is based on the cross-sectional area of each cell that is bounded by a line integral. This scheme is a simple application of Ampere's law.

Similarly, the magnetic permeability is averaged across the cell faces to correspond to the location of

the magnetic fields. A simple relation can be derived where this permeability is a geometric average of the two permeabilities of the neighboring cells such that the normal component of the magnetic induction, $\mathbf{B} = \mu\mathbf{H}$, is continuous. For example, in Figure 2 the average permeability corresponding to H_x at the face $(i + 1, j + 1/2, k + 1/2)$ uses the permeability properties of cells $(i + 1, j, k)$ and (i, j, k) , and is given by

$$\mu_{\text{avg}} = \frac{(x_{i+3/2} - x_{i+1/2})\mu_{i+1,j,k}\mu_{i,j,k}}{(x_{i+3/2} - x_{i+1})\mu_{i,j,k} + (x_{i+1} - x_{i+1/2})\mu_{i+1,j,k}}. \quad (5)$$

Notice in (1) that if the first term of the right-hand side is moved to the left side, then it contains both the derivatives as well as the unknown scattered field values, while the right-hand side consists of the

$$\hat{y}_{\text{avg}} = \frac{\hat{y}_{i,j+(1/2),k}A_{i,j+(1/2),k} + \hat{y}_{i+1,j+(1/2),k}A_{i+1,j+(1/2),k} + \hat{y}_{i,j+(1/2),k+1}A_{i,j+(1/2),k+1} + \hat{y}_{i+1,j+(1/2),k+1}A_{i+1,j+(1/2),k+1}}{A_{i,j+(1/2),k} + A_{i+1,j+(1/2),k} + A_{i,j+(1/2),k+1} + A_{i+1,j+(1/2),k+1}} \quad (4)$$

known equivalent source values. Thus we can set up a system of equations which are written in matrix form as

$$\bar{\bar{A}}\bar{\mathbf{f}} = \bar{\mathbf{s}} \quad (6)$$

where $\bar{\bar{A}}$ is an N -by- N complex symmetrical matrix containing the numerical approximation of the spatial derivatives in (1) as well as the model electrical properties, N is the number of unknown electric field values, $\bar{\mathbf{s}}$ is the primary field source vector, and $\bar{\mathbf{f}}$ is the scattered field vector we wish to solve for. To assemble a linear system of equations, we employ Dirichlet boundary conditions with the tangential electric fields at the edges of the grid set equal to zero. Unwanted reflections are avoided by either real or complex grid stretching as described below. The explicit formulation for the finite difference solution is given in Appendix B.

The solution vector is obtained by using preconditioned Krylov subspace techniques. These methods are among the fastest iterative solvers currently available for large sparse systems of equations. These methods also efficiently map to a parallel computer architecture. A brief description of the Krylov solvers employed is given in the next section.

After the scattered fields at the grid points have been determined, the fields at the receiver(s) must be calculated. The electric fields are simply calculated by bilinear interpolation of the values at the nodal points surrounding the receiver point. The magnetic fields at the receivers are calculated by first taking the numerical approximation of Faraday's law for the scattered electric fields surrounding the receiver,

$$\nabla_e \times \mathbf{E}^s = -i\omega\mu\mathbf{H}^s + (\mu - \mu_p)\mathbf{H}^p \quad (7)$$

and then interpolating to the point of interest.

The Krylov Subspace Solver

In recent years a number of Krylov subspace algorithms have been developed for handling linear systems such as ours where $\bar{\bar{A}}$ is complex symmetric and not Hermitian. These methods proceed in an iterative manner to minimize

$$\bar{\mathbf{r}} = \bar{\bar{A}}\bar{\mathbf{f}} - \bar{\mathbf{s}}. \quad (8)$$

Though a reduction in the error is not guaranteed at each iteration, and thus the convergence is generally erratic, these techniques have proven success-

Table 1. Type and Number of Major Mathematical Operations Involved in the BICG and QMR Schemes

Operation	Times Used Within the Recursion	
	BICG	QMR
Vector-matrix multiply	1	1
Vector dot product	3	4
Vector addition/subtraction	3	6
Vector constant multiply	3	9

BICG, biconjugate gradient; QMR, quasi-minimum residual.

ful in reducing the error to a predetermined level within an acceptable number of iterations. The most widely used of the techniques is the biconjugate gradient (BICG) method first employed for electromagnetic modeling by *Sarkar* [1987] and more recently for magnetotelluric modeling by *Smith* [1992]. Here we have examined both this method as well as the quasi-minimum residual (QMR) technique proposed by *Freund* [1992], who has determined that these two methods offer the best trade-off between solution accuracy and speed for matrix systems that are complex symmetrical. In order for the reader to better understand the parallelization of the FD scheme as discussed below, a general description of these routines is given here. For a more explicit description, the reader is referred to work by *Freund* [1992].

To initialize both routines, a starting solution vector $\bar{\mathbf{f}}_0$ is chosen, the residual given by (8) and a matrix-vector multiply are computed, and then various vector dot products are calculated. After this initialization process has been completed, the recursion begins. Table 1 lists the four mathematical operations which dominate the run time within the recursion, as well as the number of times each operation is employed within both of the solvers. Because the QMR scheme is more computationally intensive than the BICG method, it will take more time per iteration to complete the necessary calculations. However, as explained by *Freund* [1992], this added complexity is offset by stability.

This observation is illustrated in Figure 3, where the residual error (defined as $\|\bar{\mathbf{r}}\|^2/\|\bar{\mathbf{s}}\|^2$) for both methods is plotted as both a function of iteration number and time for one of the models discussed below. Both methods converge rather quickly up to approximately 900 iterations. However, the BICG




Figure 3. Convergence of the BICG and QMR methods for the HMD simulation in Figure 5 at 7.2 kHz. (a) Normalized residual versus run time. (b) Normalized residual versus iteration number.

method loses stability beyond this point, and the residual begins to oscillate. The net result is that it takes over twice as long to converge to the desired level when compared to the QMR method. We feel that this is most likely due to the fact that the BICG scheme experiences problems associated with round-off errors before the QMR does. Although we have found that the BICG routine will often converge 5% faster than the QMR method, we prefer the QMR scheme because of its more stable nature.

Note that when using these methods, \bar{s} has to be reconstructed for each new source, and both \bar{A} and \bar{s} for each frequency, that is, the process has to be reinitialized for each new source and frequency. However, if successive source positions and/or frequencies do not exhibit large differences from one another, then a bootstrapping technique can be performed where the previous solution vector is used as the initial guess (\bar{f}_0) for the new model. *Newman and Alumbaugh [1995]* have found that for airborne simulations this process can offer a time savings of up to 15%.

Preconditioning

In order to accelerate the convergence of the BICG, QMR, and other Krylov methods, preconditioning can be applied to the linear system. This process reduces the condition number of the system matrix by finding some approximation (\bar{M}^{-1}) to

\bar{A}^{-1} , and then multiplying through the system by this matrix [*Barrett et al., 1994*]. At this time several different preconditioning schemes have been implemented and tested including modified incomplete Cholesky methods, two different polynomial preconditioners, a block Jacobi method, and a simple diagonal or Jacobi scaling. An excellent overview of all these methods has been given by *Barrett et al. [1994]*. Of these methods, Jacobi scaling and a least squares polynomial offer the best acceleration of convergence for the system we are solving. In addition, although we have found that the polynomial preconditioner can produce up to a 25% time savings over the diagonal scaling if the optimal number of polynomial terms are employed, choosing the wrong number of terms can result in much slower convergence. Thus, because we have not fully investigated a method for choosing the optimal number of terms, Jacobi scaling was employed for all of the results presented below.

The primary reason that Jacobi scaling works as well as it does is that it only has to be applied once before the solver is called; most other preconditioners must be recalculated and applied at each iteration within the Krylov routine. The scaling involves modifying (6) such that the linear system looks like

$$[\bar{D}^{-1/2} \bar{A} \bar{D}^{-1/2}] \bar{D}^{1/2} \bar{f} = \bar{D}^{-1/2} \bar{s} \quad (9)$$

where $\bar{\bar{D}}_{ij}^{1/2} = 0$ for $i \neq j$, and $\bar{\bar{D}}_{ij}^{1/2} = (\bar{\bar{A}}_{ii})^{1/2}$ for $i = j$. The modified matrix $[\bar{\bar{D}}^{-1/2} \bar{\bar{A}} \bar{\bar{D}}^{-1/2}]$ is still complex symmetric but with unity on its main diagonal. Because the whole system is scaled before entering the solver, the scaled fields $(\bar{\bar{D}}^{1/2} \bar{\bar{r}})$ must be rescaled back to their true values after an acceptable convergence has been obtained.

Properties of the Perfectly Matched Layer Absorbing Boundary Condition

The perfectly matched layer (PML) absorbing boundary condition developed by *Berenger* [1993] was chosen over other ABCs for two reasons: (1) it is naturally parallelizable, as opposed to an integral boundary condition such as the one given by *Druskin and Knizhnerman* [1994], and (2) it leaves the matrix complex symmetrical, which is not true of many other ABCs such as the Bayliss-Turkel and Liao conditions [Chew, 1990]. In addition, the incorporation of complex grid stretching involves only slight modifications to the existing serial version of the code making it easy to implement, and unlike the time domain solution the frequency domain version does not require a doubling of memory.

Although their calculations employ the coupled modified Maxwell's equations in the time domain, *Chew and Weedon* [1994] develop theory in the frequency domain to demonstrate how lossy, non-reflecting conditions are created along the mesh boundaries. Using their type of analysis, we demonstrate that these same properties hold for the scattered electric field Helmholtz equation in Appendix A. Here we explain the physical characteristics of the complex stretching variables, how they are implemented, and, briefly, how they affect our linear system.

The complex stretching parameters (e_j and h_j for $j = x, y, z$ in (2) and (3)) are assigned a value of the form $1 + a - ib$, and the method in which they are incorporated into the discrete finite difference equation is given in Appendix B. On the internal portion of the mesh, $a = b = 0$ such that the modified Helmholtz equation reduces to the normal form. Near the edges of the mesh e_j and h_j are allowed to vary over several cells, but only in the direction that is perpendicular to the boundary. For example along the $+z$ boundary $e_x = h_x = e_y = h_y = 1$, and only e_z and h_z are allowed nonzero values of a and b .

Because we are solving an implicit rather than explicit system, we have found that in order to incorporate a given amount of loss, or attenuation, across a number of cells serving as the absorbing boundary, it is better to set a and b constant rather than gradually increasing their value toward the mesh boundaries, as suggested by *Berenger* [1993]; gradually increasing their value results in a greater number of iterations needed to achieve convergence. We believe that this property can be explained in terms of how the mesh design affects the spectral properties of $\bar{\bar{A}}$. We have found empirically that as the "aspect" ratio between the largest cell and smallest cell in the mesh increases, the convergence slows because of the larger condition number of $\bar{\bar{A}}$. Thus gradually increasing the stretching parameters outward will produce a cell along the edge of the mesh which is effectively much larger than any of the cells in the model that has employed a constant stretching. Because the smallest cell size is the same in either case, the solution of the model that employs the gradual stretching will have a larger maximum aspect ratio and thus will take longer to converge.

Currently, we do not have a rigorous method for choosing a and b . Rather they have been chosen by observing the convergence properties of the Krylov solvers for models that can be compared to analytic results. In general we have found that for frequencies less than 100 kHz, it is best to set $b = 0$ and vary only a , that is, incorporate normal grid stretching. For frequencies greater than 10 MHz, better results are obtained by setting a to 0 and varying only b . In between these frequencies a combination of a and b can be employed.

Parallel Implementation on Multiple Instruction Multiple Data (MIMD) Computers

As mentioned in the introduction, the original serial version of the code has been modified to run on massively parallel MIMD (multiple instruction multiple data) machines which can have thousands of processors. This was a necessity in order to simulate more realistic models than has previously been available. These parallel architectures are designed with a given number of processors in each direction (n_x in x , n_y in y , and n_z in z) such that the total number employed is equal to $n_x \times n_y \times n_z$.

The model is then broken up across the processor bank such that each individual processor is in charge of a 3-D subset of the model. Because each processor needs to make only the necessary calculations for this subset, and because all of the processors are making their appropriate calculations simultaneously, the solution time is reduced by a factor which is approximately equal to the number of processors employed as long as interprocessor communication is minimal.

The first step in converting the serial version of the code to a parallel version is to divide the problem up among the processors such that it is load balanced. This preprocessing step is necessary to ensure that large banks of processors are not standing idle for long periods of time while a single or small number of processors complete their calculations. Thus the problem is broken up such that each processor has as close to an equal number of unknowns as possible for which to solve.

The second issue that needs to be addressed is inputting the model. To accomplish this, we have decomposed the input data into two different sets: a global data set and a local data set. Global data are those variables that each processor needs to know, such as the source and receiver positions, the frequencies, what type of solver is being employed, and the location of the mesh nodes. These form a fairly small data set which can easily be read in by a "lead" processor and then "broadcast" to all other processors. The second type of input is the local data, or local model parameters (conductivity, dielectric permittivity, and magnetic permeability) that are assigned to each cell within the model. These form a large data set (up to hundreds of megabytes) which cannot be input like the global data because of limited amounts of local memory on each processor. However because each processor only needs a small portion of the model, the local data can be broken up into multiple files, one for each processor, which are then read in individually. Additional time savings can be incorporated by distributing these multiple files across a parallel disk system which allows several files to be read simultaneously.

After the data have been loaded, each processor constructs its own portion of the matrix $\bar{\mathbf{A}}$ and source vector $\bar{\mathbf{s}}$ and then proceeds to solve for its portion of the solution vector. However, because each iteration within the solver routines requires one matrix-vector multiply and several vector dot

products, information must be exchanged both between all of the processors as well as small subsets of processors. The dot products are fairly easy to implement as they involve (1) a local calculation in which each processor computes the dot product of its portion of the vector and (2) a global calculation in which all the local calculations are "gathered" by the lead processor, summed, and the result broadcast across the machine.

The matrix-vector multiply is more difficult to implement because it requires that each processor communicates with those "neighboring" processors that contain adjacent portions of the model. This is deduced by closely examining the finite difference stencil in Figure 1. For example, in order for the processor containing node (i, j, k) to complete its matrix-vector multiply in a given iteration, it will need to know the current values of the appropriate unknowns assigned to adjacent nodes. However, if (i, j, k) lies on a boundary such that nodes $(i, j + 1, k - 1)$ and $(i, j + 1, k)$ are assigned to a different processor, then the processor containing (i, j, k) will need to receive the updated values of those particular unknowns from the processor containing $(i, j + 1, k - 1)$ and $(i, j + 1, k)$, and vice versa. Thus the next and most difficult step in implementing the code on a parallel machine is to determine both the stencil which defines the neighboring processors that each processor must communicate with and also which unknowns will be communicated.

This is accomplished in the following manner. If we assume that each processor contains only a single node, then we can imagine it as a cubic shape enclosing node (i, j, k) as well as all other nodes in Figure 1. Careful examination then indicates that there are two types of communication that each processor needs to execute with its appropriate neighbors. The first type of communication will occur across the "faces" of the cube. For node (i, j, k) this implies communication with those nodes directly connected to it by the gray lines of the finite difference stencil, that is, nodes $(i - 1, j, k)$, $(i + 1, j, k)$, $(i, j - 1, k)$, $(i, j + 1, k)$, $(i, j, k - 1)$ and $(i, j, k + 1)$, and involves sending and receiving either two or three unknowns to each of these. The second type of communication occurs across certain "corners" of the cube and involves those nodes which are not directly connected to (i, j, k) by the stencil lines, for example, node $(i + 1, j, k - 1)$. This type of communication requires only one unknown per




Figure 4. The processor communication stencil that provides for proper message passing in the solution phase of the program. Each cube represents a neighboring processor with which a processor located at the center of the “face contributions” cluster would need to exchange information through message passing.

node being communicated each way. If we now expand the idea such that each processor cube contains a 3-D distribution of nodal points, then we can develop the processor communication stencil shown in Figure 4. Thus, if the processor we are interested in contains a certain block of the model which is represented by a cube at the center of the stencil in Figure 4, then the processors, or portions of the model that it needs to communicate with, are arranged about this specific processor in the manner illustrated.

The last issue to be addressed is the data output. Because for any given source we only need to know the results at a limited number of receiver positions, all of which may lie on the same processor, the data output is inherently nonparallel and is accomplished in the following steps: (1) Each processor determines which processor holds the portion of the model that contains the receiver position; (2) this “receiver” processor then determines if it needs any values from adjacent processors, completes the necessary point-to-point communication with those processors, and then does the necessary bilinear interpolation; and (3) the results are then sent to the lead processor which outputs them to disk.

To this point the code has been implemented on two different MIMD machines available at Sandia

National Laboratories, the 1840-processor Intel Paragon and 1024-processor NCUBE; run time characteristics for the Paragon are given below. To provide for the required message passing on these two machines we have chosen to employ the message-passing interface (MPI) [Skjellum *et al.*, 1993] rather than using machine-specific commands. This provides portability to the code as it will be able to run on any parallel machine and/or distributed network of machines on which this public domain library is available.

Demonstration of the Finite Difference Solution

To demonstrate versatility of the numerical solution, we have simulated three different models which represent measurement configurations that might be employed in the field. Two of the simulations involve one-dimensional (1-D) geometries, that is, layers, and are compared against a 1-D modeling code developed by Ki Ha Lee at Lawrence Berkeley Laboratory. This code can incorporate layers of varying thickness, conductivity, and dielectric permittivity and magnetic permeability and can calculate the response for both electric and magnetic dipole sources oriented in any direction for normal induction frequencies up into the radar range. The third model involves a comparison against a 3-D IE solution given by Newman *et al.* [1986]. In all of the cases presented the background dielectric constant and magnetic permeability were both assumed to be that of free space, and only the background conductivity varies from one model to the next. In addition, the QMR solver was assumed to have converged to an adequate error level when $\|\mathbf{r}\|^2/\|\mathbf{s}\|^2$ was found to be less than or equal to 1.0×10^{-8} . This error level is empirical and is based on extensive comparisons of the solution with other numerical solutions and scale model experiments [Alumbaugh and Newman, 1994].

Airborne Simulation

The first example crudely simulates an airborne experiment where both a vertical magnetic dipole (VMD) and x -directed horizontal magnetic dipole (HMD) are located 20 m in the air over a 100- Ωm half space (Figure 5a). The difference between this model and the one given in Figure 3a of Newman and Alumbaugh [1995] is that this incorporates a magnetic permeability of $5 \cdot \mu_0$. The sources are

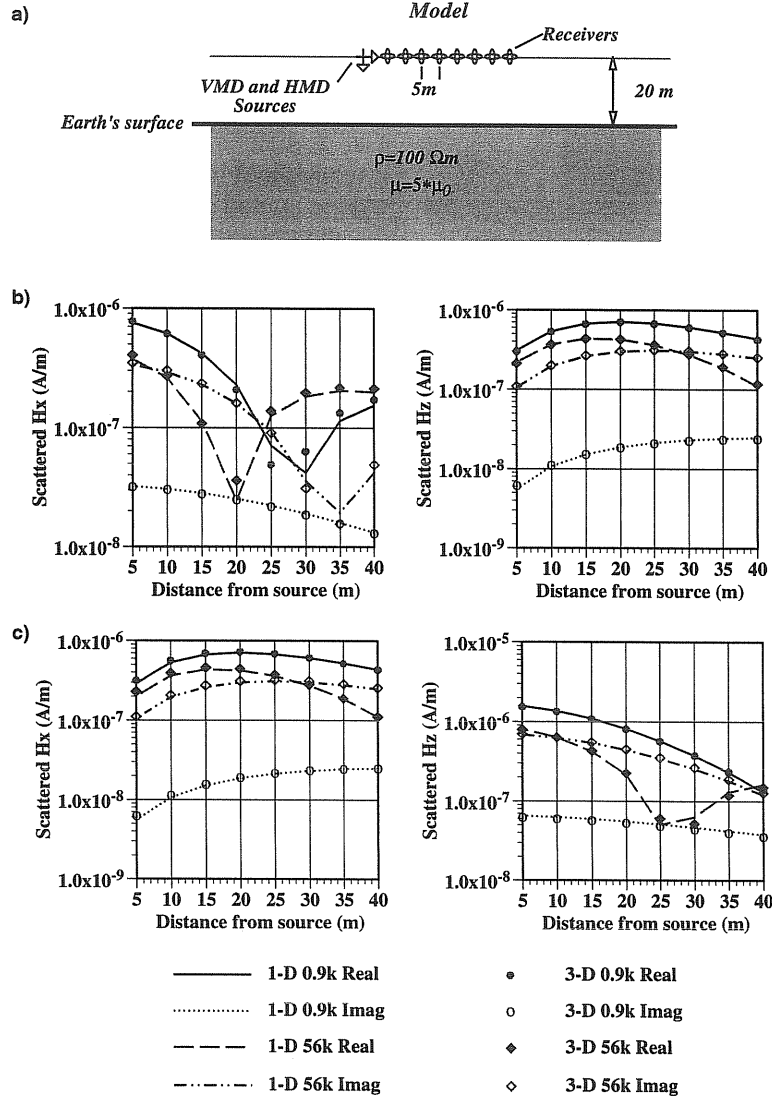


Figure 5. (a) Airborne simulation with sources and receivers above a magnetically permeable half space. The two frequencies employed are 0.9 and 56 kHz. The comparisons are between the 3-D finite difference scheme and Lee's 1-D solution. (b) Absolute values of the scattered horizontal and vertical magnetic fields for an HMD source. (c) Absolute values of the scattered horizontal and vertical magnetic fields for a VMD source.

operating at 0.9 and 56 kHz, and eight receivers are located at the same height as the sources at 5-m intervals in x . To calculate this with the 3-D finite difference code, the earth and air were divided into a $54 \times 43 \times 53$ cell grid which yields a total of 3.7×10^5 unknowns for which to solve. To avoid reflections off the mesh boundaries, normal grid stretching (i.e., $b = 0$) was employed to move them out to plus or minus 320 m in x , plus or minus 300 m in y ,

and -310 m and $+240$ m in z where the air-earth interface is at 0 m. The smallest cell size employed was 5×5 m by 2.5 m and was employed at the air-earth interface underneath the source-receiver array. The largest cell size employed was in the corners of the mesh and was $20 \times 20 \times 20$ m. A background conductivity of $\sigma = 1.0 \times 10^{-16}$ S/m was assumed to simulate the electrical properties of the air.

Table 2. Number of Iterations and Time of Solution for Examples Shown in Figure 5

Source Type Frequency, kHz	VMD		HMD	
	Iterations	Time, s	Iterations	Time, s
0.9	1373	155	1374	155
56	361	41	652	73

Figures 5b and 5c show a comparison of the magnetic fields for the model compared to the 1-D solution. These results have had the free space response removed such that only the fields which have been backscattered off the earth are plotted. Although the comparison for both sources is excellent, the results for the HMD are especially exciting as the solution for this polarization not only requires a jump to occur in the vertical magnetic field across this surface, but all vertical electric fields within the earth must go to zero. From these results it appears that the FD solution is capable of doing this. However, the solution for the HMD takes longer than that of the VMD, which indicates that this is a harder problem to solve.

The solution convergence times for the two sources using 100 processors is given in Table 2. Even though this is a very small problem that uses only a limited number of processors, the quick turn-around time clearly illustrates the advantage of using the parallel processors; running this same model on our high-end workstation (an IBM RS600-590) took over an order of magnitude longer. In addition, these results illustrate the convergence pattern that has been generally observed for frequencies below 1 MHz with the higher-frequency simulations converging quicker than the lower.

Cross-Well Simulation

The second example is a cross-well model which simulates the EM monitoring of an enhanced oil recovery process such as the one examined by *Wilt et al.* [1995]. As shown in Figure 6a, a conductive block representing either a hot water or steam stimulation is placed in a 100- Ω m layer representing an oil-bearing layer. The reservoir is bounded above and below by 20- Ω m layers which are assumed to be the resistivity of the background whole space. Two wells are located 100 m apart on each side of the block and in the first simulation a VMD source is located 40 m above the layer, while in the second case the source is within the layer. Both the

horizontal (x directed) and vertical magnetic fields are calculated at the receiver positions in the second well.

The grid size of $82 \times 62 \times 63$ cells yields a total of 9.6×10^5 unknowns. To avoid reflections off the mesh boundaries, the grid was 800 m long in x , 600 m long in y , and 600 m in depth (z). The largest cell size employed at the boundaries was $20 \times 20 \times 20$ m, while within the body a minimum cell size of $2.5 \times 2.5 \times 2$ m was employed. In this case 252 processors were employed on the Paragon. The total run time for both sources and frequencies was 26 min, with the 0.1-kHz simulations taking approximately 5000 iterations to converge, and the 10-kHz runs taking approximately 700 iterations.

To check the finite difference results we calculated the response of the block in a layered half space using the integral equation scheme described by *Newman et al.* [1986]. To discretize the block, cubic cells 5 m on a side were employed, which resulted in a model that took a few minutes to run for each frequency on our IBM RS6000 workstation. The comparison for the source above the layer is given in Figure 6b and that for the source within the layer in Figure 6c. Notice that in all cases the comparison between the two solutions is excellent. It is especially promising that even when the source is located within the layer and thus embedded in a zone of anomalous conductivity, the comparison is good with normal discretization about the transmitter.

Figure 7 shows results for the same exact geologic model as given above with the VMD source replaced by a vertical electric dipole (VED) source. Thus here we are comparing the horizontal and vertical electric fields for the two solutions. Figure 7b shows that for the source located above the layer the results are in excellent agreement. Even the rapid discontinuity at the layer boundaries in the z component of the field is mapped, although it is done so in a slightly smoother fashion. This smoother transition for the FD calculations is due to the manner in which the conductivity is averaged between adjacent cells and to the manner in which the bilinear interpolation scheme averages adjacent data. The solution time for the VED was slightly longer than that of the VMD, with the 0.1-kHz simulation taking approximately 6000 iterations and the 10-kHz run taking 780 iterations for the sources located above the layer.

The comparisons shown in Figure 7c with the

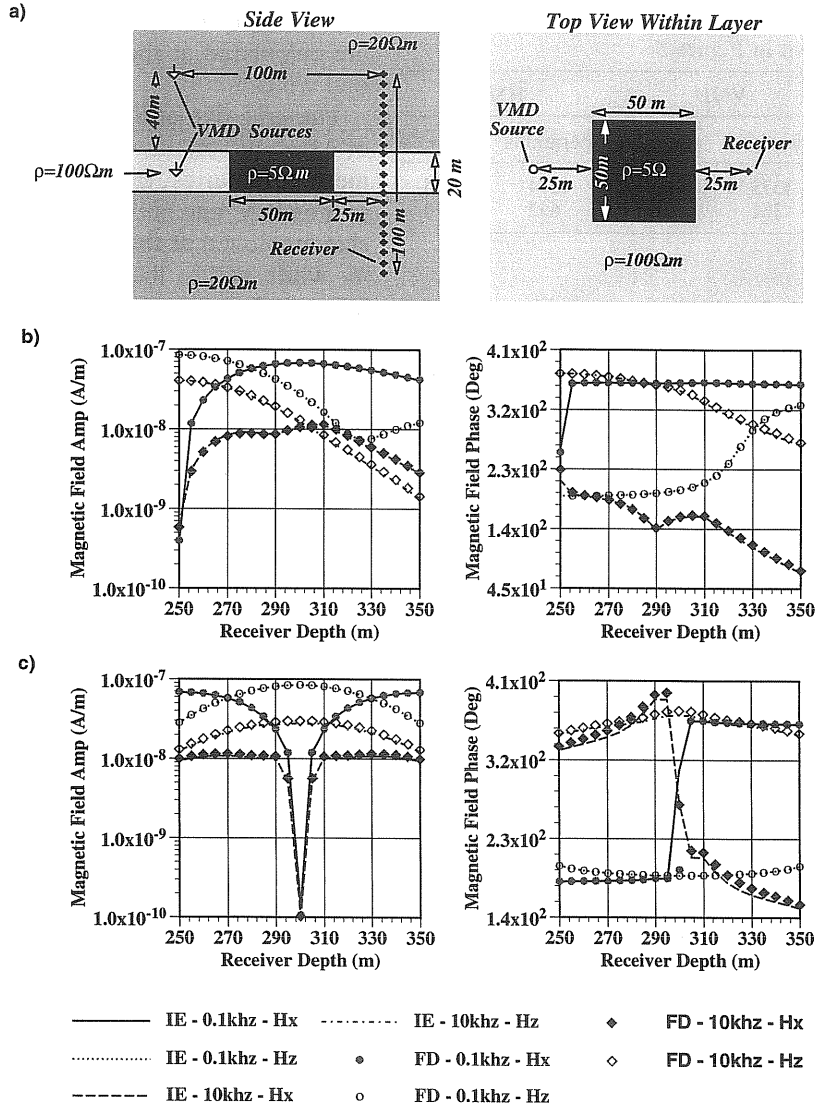


Figure 6. (a) Cross-well model with VMD source located 40 m above and within a resistive layer. The receiver depths range from 250 to 350 m, and the frequencies employed are 0.1 and 10 kHz. The comparisons are between the 3-D finite difference scheme and the integral equation solution of Newman *et al.* [1986]. (b) Horizontal and vertical magnetic field results for the source above the layer. (c) Horizontal and vertical magnetic field results for the source within the layer.

VED source within the layer, though very good, were much more difficult to obtain compared to the three other cross-well examples given above. The most serious problem presented to the FD solution was that the VED source was located in a zone of anomalous conductivity. When the background whole space was assumed to be 20 Ωm , the amplitudes were off by almost an order of magnitude. To get the good comparison shown in Figure 7c, a

background of 100 Ωm , that is, equal to that of the layer, was assumed. This result demonstrates that when the source is located within or near a region of anomalous electromagnetic properties, caution must be exercised depending on the source type and polarization.

It must also be mentioned here that although the results in Figure 7c compare very well, the comparison is not nearly as good as it was for the other

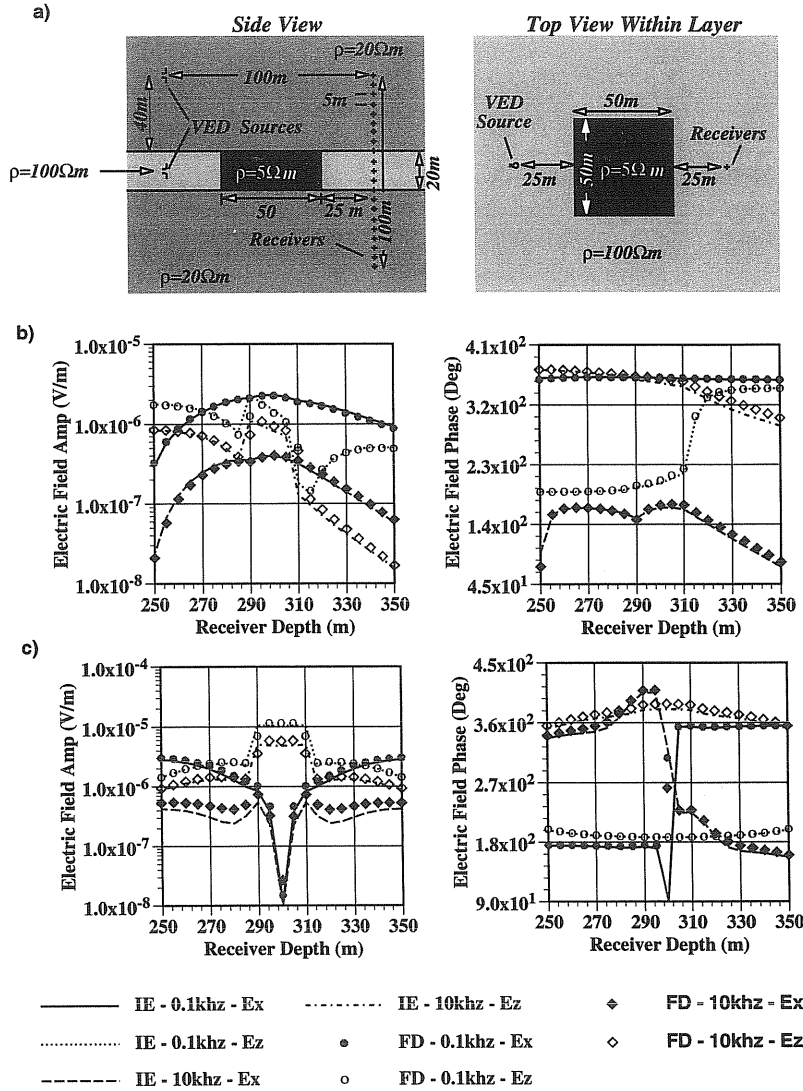


Figure 7. (a) Cross-well model with VED source located 40 m above and within a resistive layer. The receiver depths range from 250 to 350 m, and the frequencies employed are 0.1 and 10 khz. The comparisons are between the 3-D finite difference scheme and the integral equation solution of Newman *et al.* [1986]. (b) Horizontal and vertical electric field results for the source above the layer. (c) Horizontal and vertical electric field results for the source within the layer.

cases, especially in amplitude. However, we believe that this is not the fault of the FD scheme but rather a discretization problem with the IE solution. We have arrived at this conclusion using the simulation shown in Figure 8. Changing the discretization within the FD simulation had almost negligible results. However, the results are very different depending on whether a 5-m or 2.5-m cubic cell dimension is employed to discretize the body within the IE scheme. Unfortunately, we could not further

decrease the cell size because due to memory limitations of our IBM workstation; further decreasing the cell size to 1.25 m^3 would have required eight times more memory. Nevertheless, because the IE results calculated with the 2.5-m cubic cells compare to the FD results much better than those results calculated using a 5-m cubic cell size, we believe that a closer comparison could be obtained if a finer discretization was employed within the IE scheme.

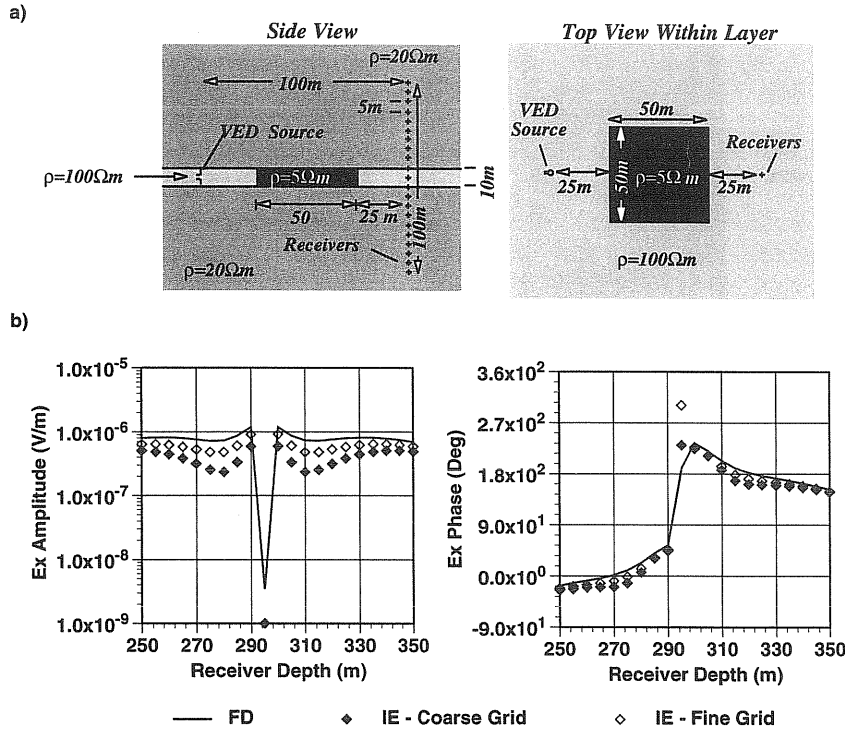


Figure 8. (a) Cross-well model with VED source located within a resistive layer. The receiver depths range from 250 to 350 m, and the frequency employed is 10 kHz. The comparisons are between the 3-D finite difference scheme and the integral equation solution of *Newman et al.* [1986] using two different discretizations. (b) Horizontal electric field results showing how using a finer discretization within the IE scheme improves the fit to the FD results.

High-Frequency Simulation and the Need for Absorbing Boundary Conditions

In order to help in the design of the “VETEM” (very early time electromagnetic) geophysical prospecting system [Pellerin *et al.*, 1995] the code must be able to accurately calculate the response for frequencies ranging from 500 kHz to 50 MHz. To illustrate the accuracy of the code at these frequencies and the need for absorbing boundary conditions, the model shown in Figure 9a has been employed. This example was designed to simulate a test site at the Colorado School of Mines, where a prototype of the VETEM system known as the “high-frequency sounder” [Stewart *et al.*, 1994] was first tested. The model is particularly difficult to simulate because of two conflicting conditions that are imposed by the material properties: (1) The wavelength in the second layer at 28.5 MHz is approximately 1.6 m, which requires a maximum cell dimension of 0.16 m to avoid grid dispersion [Chew, 1990, p. 244], and (2) the skin depth in the

first layer at that same frequency is 17.8 m, which requires the boundaries to be placed very far away to avoid reflections off the grid. The small cell size coupled with the large distance to the boundaries produces a very large mesh if no absorbing boundary conditions are employed. In addition, as demonstrated below, simply enlarging the cells at the boundaries as we have done in the previous examples does not work and actually can make conditions worse due to grid dispersion. Thus absorbing boundary conditions are needed to solve the problem.

Throughout most of the following examples, a $120 \times 120 \times 120$ cell mesh was employed with a constant cell size of 0.15 m in the x and y directions. This places the total distance across the mesh at 18 m. In z the maximum cell size was also 0.15 m, with a minimum cell size of 0.13 m to accommodate the layer thicknesses. Note that this mesh produces a total of five million unknowns. The VMD source was placed at the center of the mesh in x and y , that

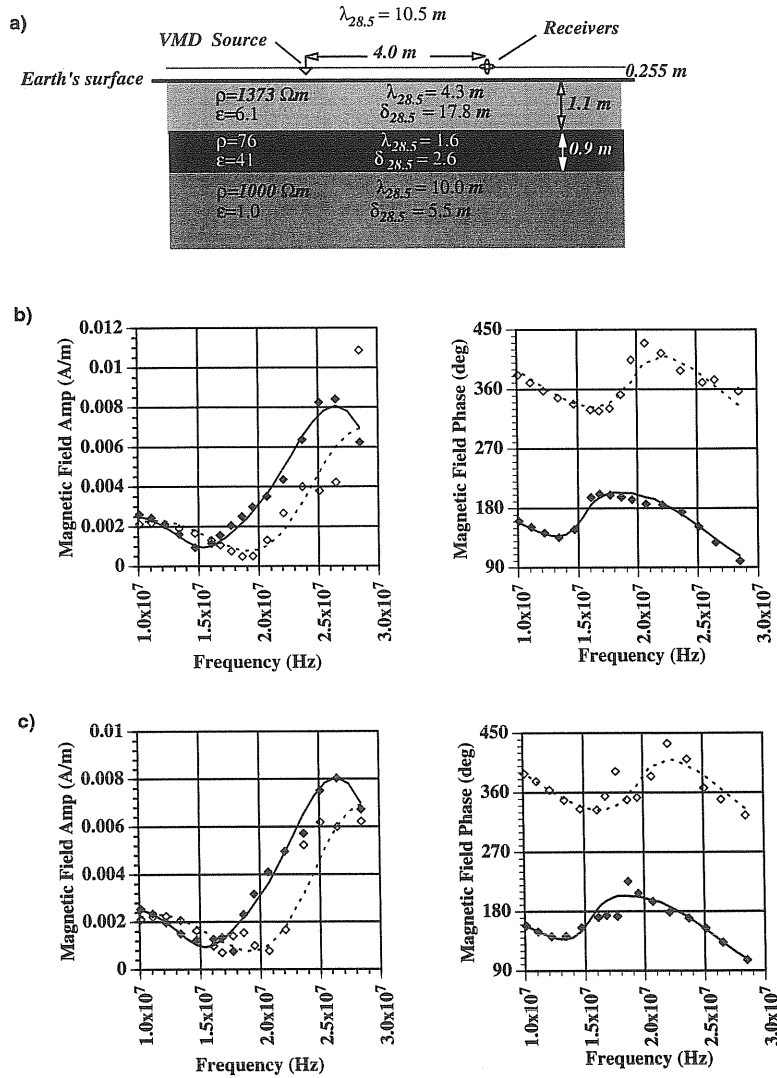


Figure 9. (a) Colorado School of Mines 1-D model. The magnetic permeability in all three layers has been set to that of free space. Unless otherwise noted a $120 \times 120 \times 120$ cell mesh was employed with a maximum cell dimension of 0.15 m. The comparisons are between the 3-D finite difference scheme (symbols) and Lee's 1-D solution (lines). (b) Horizontal and vertical magnetic field results obtained with no grid stretching. (c) Horizontal and vertical magnetic field results obtained with real grid stretching ($a = 1.0$, $b = 0.0$) for 25 cells along each boundary. (d) Horizontal and vertical magnetic field results obtained with complex grid stretching ($a = 0.0$, $b = 0.6$) for 25 cells along each boundary. (e) Horizontal and vertical magnetic field results obtained on a $72 \times 72 \times 72$ cell mesh with increased complex grid stretching ($a = 0.0$, $b = 2.0$) for 10 cells along each boundary.

is, 9 m from each boundary, and a background conductivity of $\sigma = 1.0 \times 10^{-16}$ S/m was assumed.

In Figure 9b the horizontal and vertical magnetic fields calculated with the 3-D code and no absorption have been plotted against Lee's 1-D solution. It is immediately evident that the 3-D solution begins to break down at about 15 MHz, and we can assume

that this is due to reflections of the mesh boundaries contaminating the solution. Doubling the size of cells along the mesh boundaries, that is, using normal grid stretching, does not help matters. Figure 9c shows that poor results occur when a real stretching parameter $a = 1.0$ is employed along 25 cells of each boundary. However, when complex

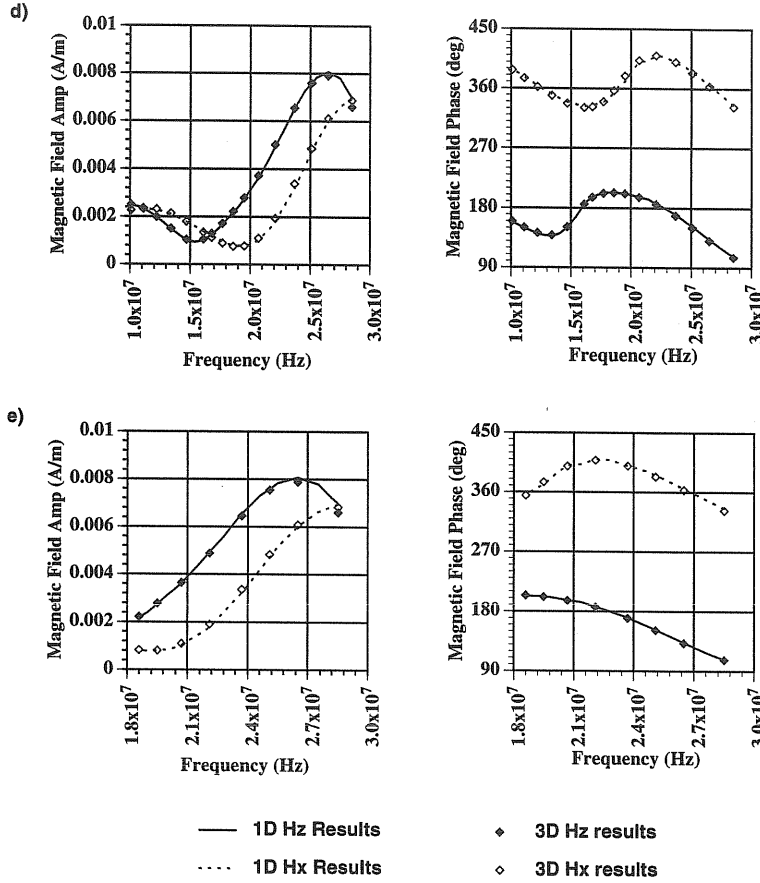


Figure 9. (continued)

grid stretching is employed the results are much better. Figure 9d shows that when a stretching parameter of $b = 0.6$ is employed along 25 cells of each boundary, the fields calculated with the 3-D solution match those of 1-D solution almost exactly.

In Figure 9e we demonstrate how the absorbing boundary conditions can be employed to shrink the size of the mesh. In this case a $72 \times 72 \times 72$ mesh was employed, and again a maximum cell size of $0.15 \times 0.15 \times 0.15$ m was used. This small mesh size places the boundaries only 10.8 m apart. To calculate the results, a complex grid stretching parameter of $b = 2.0$ was employed along 10 cells of each boundary. Notice that the 3-D calculations again agree almost exactly with the 1-D solution. This example fully illustrates the utility of these absorbing boundary conditions as not only do they allow one to accurately model high frequency where propagation rather than diffusion dominates,

but they also allow the mesh size to be significantly reduced, which results in much quicker run times.

Because the rate of convergence of the Krylov solvers is dependent on the spectral properties of the linear system, an analysis of how the PML boundary condition affects the system can be accomplished by observing how different stretching parameters alter this convergence. Thus in Figure 10 the number of iterations for convergence have been plotted for those results shown in Figures 9b through 9d. Notice that at low frequencies, the unstretched system converges quickly. However, as the frequency increases, the number of iterations needed for convergence also increases. The rapid increase in solution time is even more dramatic for system that has had real grid stretching applied. In contrast, the system with the complex grid stretching behaves much differently. Although it takes longer to converge when compared to either of the




Figure 10. The solution times of the QMR method for different frequencies and the different grid stretching parameters employed for Figures 9b through 9d.

other methods at lower frequencies, it does not experience the rapid increase in solution time as the frequency increases. This indicates that unlike the other two methods of terminating the mesh, the spectral properties of the system that employs complex grid stretching is almost independent of frequency, at least for frequencies greater than 10 MHz.

Run Time Properties of the Solution on the Intel Paragon

To fully demonstrate the power of the massively parallel (MP) code and to demonstrate some of the questions that must be answered when using these machines, the solution time has been plotted against the number of processors employed for one of the VETEM simulations examined above. Figure 11 shows the solution times versus the number of processors employed as well as the flop rate at which the Paragon is operating. The large decrease

in run time with an increasing number of processors up to 512 indicates that the processors are spending the majority of their time performing calculations rather than communicating. This corresponds to solving for 10,000 to 24,000 unknowns per processor. Medium processor efficiency is obtained when between 512 and 1000 processors are employed, which corresponds to 10,000 to 3,000 unknowns being assigned to each processor. The relatively small decrease in run time with increasing number of processors over 1000 indicates that if fewer than 3000 unknowns are being solved for on each processor, the code is inefficient due to message passing. Thus, if we wish to use the machine most efficiently, we should employ the minimum number of processors that the local memory will allow. We could then run several jobs simultaneously such that the efficiency increases proportionally to the number of jobs. On the other hand, if we desire as quick a turn-around time as possible for a single




Figure 11. The solution time and flop rate versus the number of processors employed on the Intel Paragon for the 10.1-Mhz simulation plotted in Figure 9d.

computation, then we would want to operate near the right end of the curve.

The second topic to be addressed is the maximum size of the model that we can simulate on the Paragon. Because we must assign a certain number of processors to each of the x , y , and z directions, we have found most complete use of all 1840 processors occurs when 14 are assigned to one direction, 13 to the second, and 10 to the third for a total of 1820 processors. Each node on the Paragon is equipped with 16 Mbytes of local memory. If we store the matrix \bar{A} in single precision to make maximum use of the memory, the maximum problem that a single processor can accommodate is $20 \times 20 \times 20$ cells, or 24,000 unknowns. Thus the maximum problem size is $280 \times 260 \times 200$ cells for a total of 43.6 million unknowns. However, if for accuracy we wish to store the matrix in double precision, then the maximum number of cells each processor can accommodate is $15 \times 15 \times 15$ cells, which is equivalent to 10,125 unknowns. In this case the maximum model size is $210 \times 195 \times 150$, which yields a total of 18.4 million unknowns. Note that the example running on 216 processors in Figure 11 assumes that each processor is running at maximum capacity. Taking the flop rate for this example (1.75 Gflops/s) and scaling it upward for the maximum number of processors that we can access at one time (1820) yields a theoretical maximum flop rate of 14.7 Gflops/s.

Conclusions and Discussion

In this paper we have presented a scheme to solve for the frequency domain electromagnetic response of a 3-D earth over a wide band of frequencies using massively parallel computers. The problems associated with porting the serial version of the scheme to a parallel machine have been outlined, and a variety of comparisons have been demonstrated to prove the validity of the code. Implementing the code on the 1840-processor Intel Paragon has demonstrated a decrease in computing time of over 2 orders of magnitude when compared to a high-end IBM workstation and a similar magnitude increase in the maximum model size that can be simulated. In addition a maximum theoretical flop rate of 14.9 Gflops/s has been established.

Currently we are using the scheme in a variety of projects, for example to assist in the design of

geophysical instruments [Pellerin *et al.*, 1995] as well as simulating airborne EM surveys [Alumbaugh and Newman, 1995; Newman and Alumbaugh, 1995]. The simulations that we are running for these projects would have been impossible prior to the parallel implementation due to the size of the models and/or the number of frequencies and sources involved. We believe, however, that are still several areas where the scheme can be improved. Some topics of ongoing research include the following:

1. Better preconditioners. Some techniques being considered are multigrid preconditioners and methods to separately treat the real and imaginary components of the matrix system.

2. A thorough study of the grid-stretching parameters and the development of empirical and/or analytical methods to determine the optimal stretching parameter(s) for any model.

3. A better method of interpolating the calculated fields to the receiver when it is located near a discontinuity in material properties. Currently, the interpolation scheme does not take into account the "jump" condition across the boundaries. To correctly calculate the fields at these boundaries, the normal electric current and magnetic induction, rather than just the fields themselves, must be interpolated.

4. A scheme to accelerate the convergence for very low frequency simulations where channeling currents dominate. This is needed in order to simulate natural field measurements as well as extend the frequency band down below 100 Hz. Smith [1992] has found that a static correction can be incorporated to accommodate this.

5. Better methods of dealing with the air-earth interface. We have found that this interface tremendously complicates the numerical problem, especially when electric dipole sources are employed on the surface.

A last topic that should be addressed is a comment regarding the computational differences between the frequency and time domain formulations. Because the frequency domain requires complex math to be employed, because extra precision is needed to avoid rounding errors in the Krylov methods, and because the matrix \bar{A} must be stored, the frequency domain solution is significantly more memory intensive than a finite difference time domain (FDTD) time-stepping code. In addition, to get a full time solution in some instances may only

require a few more time steps than the number of iterations required to calculate the response at a single frequency. Thus it would initially appear that a frequency domain is a waste of resources. However, for the frequency domain nonlinear inversion problem in which this code will ultimately be employed, we may only require a few frequencies over several decades. In addition, the data produced by many commercial EM geophysical systems is logarithmically rather than linearly sampled in frequency. This type of simulation can be very expensive to compute using an FDTD code and then Fourier transforming because of the large number of time steps needed to accurately calculate the low-frequency components as well as the high-frequency ones, especially if we are operating below 1 MHz, where diffusion dominates. Furthermore, in the future we will be examining the effects of dispersion which is much easier to implement in the frequency domain. Thus although the scheme presented here is more memory intensive compared to standard FDTD schemes, it is extremely valuable in simulating the earth's response at discrete frequencies.

Appendix A: Proof of the PML Absorbing Boundary Condition Using Complex Grid Stretching

The perfectly matched layer concept proposed in 2-D by *Berenger* [1993] and 3-D by *Katz et al.* [1994] was originally developed for time domain simulation of Maxwell's equations. Here we demonstrate that this method is valid for the 3-D frequency domain Helmholtz equation for the scattered electric fields using the method of *Chew and Weedon* [1994]. First we simplify (1) by assuming that we are at a boundary far away from any zones of anomalous electrical properties such that it can be written

$$\nabla_h \times \nabla_e \times \mathbf{E}^s + i\omega\mu_p(\sigma_p + i\omega\varepsilon_p)\mathbf{E}^s = 0. \quad (\text{A1})$$

Because we are far away from any anomalous zones, a possible plane wave solution to (A1) along this particular boundary is given by

$$\mathbf{E}^s = \mathbf{E}_0^s e^{i\mathbf{k} \cdot \mathbf{r}} \quad (\text{A2})$$

where $\mathbf{k} = k_x \hat{\mathbf{i}} + k_y \hat{\mathbf{j}} + k_z \hat{\mathbf{k}}$ and $\mathbf{r} = x\hat{\mathbf{i}} + y\hat{\mathbf{j}} + z\hat{\mathbf{k}}$. Because $\nabla \times \mathbf{E}_0^s e^{i\mathbf{k} \cdot \mathbf{r}} = i\mathbf{k} \times \mathbf{E}_0^s e^{i\mathbf{k} \cdot \mathbf{r}}$, it is easy to show that when (A2) is substituted into (A1), the resulting expression has the form

$$-\mathbf{k}_h \times \mathbf{k}_e \times \mathbf{E}^s + i\omega\mu_p(\sigma_p + i\omega\varepsilon_p)\mathbf{E}^s = 0 \quad (\text{A3})$$

where

$$\mathbf{k}_e = \frac{k_x}{e_x} \hat{\mathbf{i}} + \frac{k_y}{e_y} \hat{\mathbf{j}} + \frac{k_z}{e_z} \hat{\mathbf{k}} \quad (\text{A4})$$

and

$$\mathbf{k}_h = \frac{k_x}{h_x} \hat{\mathbf{i}} + \frac{k_y}{h_y} \hat{\mathbf{j}} + \frac{k_z}{h_z} \hat{\mathbf{k}}. \quad (\text{A5})$$

Using a vector identity, the left hand term in (A3) can be expanded to yield

$$(\mathbf{k}_h \cdot \mathbf{k}_e)\mathbf{E}^s - \mathbf{k}_e(\mathbf{k}_h \cdot \mathbf{E}^s) + i\omega\mu_p(\sigma_p + i\omega\varepsilon_p)\mathbf{E}^s = 0. \quad (\text{A6})$$

Because we are in a homogenous region absent of any free "secondary" charge,

$$\nabla_h \cdot \mathbf{E}^s = \mathbf{k}_h \cdot \mathbf{E}^s = 0 \quad (\text{A7})$$

and thus we are left with

$$(\mathbf{k}_h \cdot \mathbf{k}_e)\mathbf{E}^s + i\omega\mu_p(\sigma_p + i\omega\varepsilon_p)\mathbf{E}^s = 0 \quad (\text{A8})$$

or

$$\mathbf{k}_h \cdot \mathbf{k}_e = \frac{1}{e_x h_x} k_x^2 + \frac{1}{e_y h_y} k_y^2 + \frac{1}{e_z h_z} k_z^2 = \kappa^2 \quad (\text{A9})$$

where $\kappa^2 = -i\omega\mu_p(\sigma_p + i\omega\varepsilon_p)$.

Let us now assume that the plane wave is obliquely incident on an interface at $z = c$ where c is constant. *Chew and Weedon* [1994] show that the solution to (A9) is that of a 3-D ellipsoid which is satisfied by

$$k_x = \kappa \sqrt{e_x h_x} \sin \theta \cos \phi \quad (\text{A10})$$

$$k_y = \kappa \sqrt{e_y h_y} \sin \theta \sin \phi \quad (\text{A11})$$

and

$$k_z = \kappa \sqrt{e_z h_z} \cos \theta. \quad (\text{A12})$$

In addition, they find the reflection coefficients for the TE and TM modes at the boundary to be

$$R^{\text{TE}} = \frac{k_{1z} e_{2z} \mu_2 - k_{2z} e_{1z} \mu_1}{k_{1z} e_{2z} \mu_2 + k_{2z} e_{1z} \mu_1} \quad (\text{A13})$$

and

$$R^{\text{TM}} = \frac{k_{1z}h_{2z}\hat{y}_2 - k_{2z}h_{1z}\hat{y}_1}{k_{1z}h_{2z}\hat{y}_2 + k_{2z}h_{1z}\hat{y}_1} \quad (\text{A14})$$

where 1 represents the properties of medium the incident wave is traveling through, 2 designates the medium it will be transmitted to, and $\hat{y}_i = \sigma_i + i\omega\epsilon_i$.

Phase matching will occur if $k_{1x} = k_{2x}$ and $k_{1y} = k_{2y}$. To accomplish this we first set the material properties of the two media to be identical ($\kappa_1 = \kappa_2$) and then choose $h_x = e_x$ and $h_y = e_y$. If we now let $e_{1x} = e_{2x} = e_{1y} = e_{2y} = 1$, and furthermore set $\theta_1 = \theta_2$ and $\phi_1 = \phi_2$, then the two reflection coefficients in (A13) and (A14) are zero and no reflections are generated at the interface. However, by making e_{2z} complex, we provide additional loss in k_{2z} , which causes the wave to more rapidly attenuate in medium 2 than it would otherwise.

It must be mentioned that three assumptions have been made in this analysis which cannot be incorporated into the 3-D FD modeling scheme. The first assumption is that e_x , e_y , h_x , and h_y do not vary along the “z” interface. In the corners of the mesh these values are also varying to incorporate absorption along the x and y interfaces, and thus perfect matching cannot occur in these locations, and reflections will be generated. However, we have not experienced any serious problems with regards to this phenomenon. The second assumption is that we have made here is that the interface is located far away from any regions of anomalous electrical properties. Nevertheless, as the results in Figures 10d and 10e indicate, the PML is valid even when the stretching occurs within these regions, for instance, at the mesh boundary located at the bottom of the model shown in Figure 10a. Finally, it is assumed that $h_j = e_j$ for $j = x, y$, and z . As shown in Appendix B, h_j is actually a weighted average of the e_j values assigned to two adjacent cells where the weighting depends on the cell dimensions. However by using both a constant value of e_j and a constant cell size throughout the PML region, any problems with this assumption can be avoided.

Appendix B: Discrete Form of the Modified Electric Field Helmholtz Equation for a Staggered Grid

We start with the modified Helmholtz equation as given in (1). First let us expand out the two first-order curl operators, that is,

$$\nabla_e \times \mathbf{E}^s = \left(\frac{1}{e_y} \frac{\partial E_z^s}{\partial y} - \frac{1}{e_z} \frac{\partial E_y^s}{\partial z} \right) \hat{\mathbf{i}} + \left(\frac{1}{e_z} \frac{\partial E_x^s}{\partial z} - \frac{1}{e_x} \frac{\partial E_z^s}{\partial x} \right) \hat{\mathbf{j}} + \left(\frac{1}{e_x} \frac{\partial E_y^s}{\partial x} - \frac{1}{e_y} \frac{\partial E_x^s}{\partial y} \right) \hat{\mathbf{k}} \quad (\text{B1})$$

and

$$\begin{aligned} \nabla_h \times \left[\frac{(\mu - \mu_p)}{\mu} \mathbf{H}^p \right] &= \left(\frac{1}{h_y} \frac{\partial}{\partial y} \frac{(\mu_z - \mu_p) H_z^p}{\mu_z} - \frac{1}{h_z} \frac{\partial}{\partial z} \frac{(\mu_y - \mu_p) H_y^p}{\mu_y} \right) \hat{\mathbf{i}} \\ &+ \left(\frac{1}{h_z} \frac{\partial}{\partial z} \frac{(\mu_x - \mu_p) H_x^p}{\mu_x} - \frac{1}{h_x} \frac{\partial}{\partial x} \frac{(\mu_z - \mu_p) H_z^p}{\mu_z} \right) \hat{\mathbf{j}} \\ &+ \left(\frac{1}{h_x} \frac{\partial}{\partial x} \frac{(\mu_y - \mu_p) H_y^p}{\mu_y} - \frac{1}{h_y} \frac{\partial}{\partial y} \frac{(\mu_x - \mu_p) H_x^p}{\mu_x} \right) \hat{\mathbf{k}}. \end{aligned} \quad (\text{B2})$$

In this expression, μ_w for $w = x, y$, and z represents the magnetic permeability that is averaged across the face of two cells in the w th direction. Next, expanding the second-order curl on the left side of (1) we find that

$$\begin{aligned} \nabla_h \times \left[\frac{\mu_p}{\mu} \nabla_e \times \mathbf{E}^s \right] &= \left(\frac{1}{h_y} \frac{\partial}{\partial y} \left(\frac{\mu_p}{\mu_z e_x} \frac{\partial E_y^s}{\partial x} \right) - \frac{1}{h_y} \frac{\partial}{\partial y} \left(\frac{\mu_p}{\mu_z e_y} \frac{\partial E_x^s}{\partial y} \right) \right. \\ &\quad \left. - \frac{1}{h_z} \frac{\partial}{\partial z} \left(\frac{\mu_p}{\mu_y e_z} \frac{\partial E_x^s}{\partial z} \right) + \frac{1}{h_z} \frac{\partial}{\partial z} \left(\frac{\mu_p}{\mu_y e_x} \frac{\partial E_z^s}{\partial x} \right) \right] \hat{\mathbf{i}} \\ &+ \left[\frac{1}{h_z} \frac{\partial}{\partial z} \left(\frac{\mu_p}{\mu_x e_y} \frac{\partial E_z^s}{\partial y} \right) + \frac{1}{h_z} \frac{\partial}{\partial z} \left(\frac{\mu_p}{\mu_x e_z} \frac{\partial E_y^s}{\partial z} \right) \right. \\ &\quad \left. - \frac{1}{h_x} \frac{\partial}{\partial x} \left(\frac{\mu_p}{\mu_z e_x} \frac{\partial E_y^s}{\partial x} \right) + \frac{1}{h_x} \frac{\partial}{\partial x} \left(\frac{\mu_p}{\mu_z e_y} \frac{\partial E_x^s}{\partial y} \right) \right] \hat{\mathbf{j}} \\ &+ \left[\frac{1}{h_x} \frac{\partial}{\partial x} \left(\frac{\mu_p}{\mu_y e_z} \frac{\partial E_x^s}{\partial z} \right) + \frac{1}{h_x} \frac{\partial}{\partial x} \left(\frac{\mu_p}{\mu_y e_x} \frac{\partial E_z^s}{\partial x} \right) \right. \\ &\quad \left. - \frac{1}{h_y} \frac{\partial}{\partial y} \left(\frac{\mu_p}{\mu_x e_y} \frac{\partial E_z^s}{\partial y} \right) + \frac{1}{h_y} \frac{\partial}{\partial y} \left(\frac{\mu_p}{\mu_x e_z} \frac{\partial E_y^s}{\partial z} \right) \right] \hat{\mathbf{k}}. \end{aligned} \quad (\text{B3})$$

Now let us examine the parts of the Helmholtz equation corresponding to $\hat{\mathbf{i}}$, $\hat{\mathbf{j}}$, and $\hat{\mathbf{k}}$ separately

since these are the three equations that we are going to be solving at each node. For the \hat{i} component we have

$$\begin{aligned} & \frac{1}{h_y} \frac{\partial}{\partial y} \left(\frac{\mu_p}{\mu_z e_x} \frac{\partial E y^s}{\partial x} \right) - \frac{1}{h_y} \frac{\partial}{\partial y} \left(\frac{\mu_p}{\mu_z e_y} \frac{\partial E x^s}{\partial y} \right) \\ & - \frac{\mu_p}{h_z} \frac{\partial}{\partial z} \left(\frac{1}{\mu_y e_z} \frac{\partial E x^s}{\partial z} \right) \\ & + \frac{1}{h_z} \frac{\partial}{\partial z} \left(\frac{\mu_p}{\mu_y e_x} \frac{\partial E z^s}{\partial x} \right) + i\omega \mu_p (\sigma + i\omega \epsilon) E x^s \\ & = -i\omega \mu_p (\sigma - \sigma_p + i\omega (\epsilon - \epsilon_p)) E x^p \\ & - \frac{i\omega \mu_p}{h_y} \frac{\partial}{\partial y} \frac{(\mu_z - \mu_p) H z^p}{\mu_z} \\ & - \frac{i\omega \mu_p}{h_z} \frac{\partial}{\partial z} \frac{(\mu_y - \mu_p) H y^p}{\mu_y}. \end{aligned} \quad (B4)$$

Approximating (B4) with finite differences yields

$$\begin{aligned} & \left\{ \frac{\mu_p}{\mu_{z_{i+(1/2),j+(1/2),k}}} \left[\frac{1}{(e_{xi} \Delta x_i)} (E y_{i+1,j+(1/2),k}^s - E y_{i,j+(1/2),k}^s) \right. \right. \\ & \left. \left. - \frac{1}{(e_{yi} \Delta y_j)} (E x_{i+(1/2),j+1,k}^s - E x_{i+(1/2),j,k}^s) \right] \right. \\ & \left. - \frac{\mu_p}{\mu_{z_{i+(1/2),j-(1/2),k}}} \left[\frac{1}{(e_{xi} \Delta x_i)} (E y_{i+1,j-(1/2),k}^s - E y_{i,j-(1/2),k}^s) \right. \right. \\ & \left. \left. - \frac{1}{(e_{yj-1} \Delta y_{j-1})} (E x_{i+(1/2),j,k}^s - E x_{i+(1/2),j-1,k}^s) \right] \right\} \\ & \cdot \frac{1}{(h_{yj} \Delta y_j)} + \left\{ \frac{\mu_p}{\mu_{y_{i+(1/2),j,k+(1/2)}}} \left[\frac{1}{(e_{xi} \Delta x_i)} (E z_{i+1,j,k+(1/2)}^s \right. \right. \\ & \left. \left. - E z_{i,j,k+(1/2)}^s) - \frac{1}{(e_{zk} \Delta z_k)} (E x_{i+(1/2),j,k+1}^s - E x_{i+(1/2),j,k}^s) \right] \right. \\ & \left. - \frac{\mu_p}{\mu_{y_{i+(1/2),j,k-(1/2)}}} \left[\frac{1}{(e_{xi} \Delta x_i)} (E z_{i+1,j,k-(1/2)}^s - E z_{i,j,k-(1/2)}^s) \right. \right. \\ & \left. \left. - \frac{1}{(e_{zk-1} \Delta z_{k-1})} (E x_{i+(1/2),j,k}^s - E x_{i+(1/2),j,k-1}^s) \right] \right\} \\ & \cdot \frac{1}{(h_{zk} \Delta z_k)} + i\omega \mu_p \hat{y}_p E x_{i+(1/2),j,k}^s \end{aligned}$$

$$\begin{aligned} & = -i\omega \mu_p (\hat{y}_{i+(1/2),j,k} - \hat{y}_p) E x_{i+(1/2),j,k}^p \\ & - i\omega \mu_p \left\{ \left[\frac{(\mu_{y_{i+(1/2),j,k+(1/2)}} - \mu_p)}{\mu_{y_{i+(1/2),j,k+(1/2)}}} H y_{i+(1/2),j,k+(1/2)}^p \right. \right. \\ & \left. \left. - \frac{(\mu_{y_{i+(1/2),j,k-(1/2)}} - \mu_p)}{\mu_{y_{i+(1/2),j,k-(1/2)}}} H y_{i+(1/2),j,k-(1/2)}^p \right] \frac{1}{(h_{zk} \Delta z_k)} \right. \\ & \left. - \left[\frac{(\mu_{z_{i+(1/2),j+(1/2),k}} - \mu_p)}{\mu_{z_{i+(1/2),j+(1/2),k}}} H z_{i+(1/2),j+(1/2),k}^p \right. \right. \\ & \left. \left. - \frac{(\mu_{z_{i+(1/2),j-(1/2),k}} - \mu_p)}{\mu_{z_{i+(1/2),j-(1/2),k}}} H z_{i+(1/2),j-(1/2),k}^p \right] \frac{1}{(h_{yj} \Delta y_j)} \right\} \end{aligned} \quad (B5)$$

where now $\hat{y} = \sigma + i\omega \epsilon$. In this expression Δw_l for $w = x, y, z$ and $l = i, j, k$ represents the width of the l th cell in the w th direction. Similarly $\Delta \bar{w}_l$ is the distance in the w th direction between the centers of cells l and $l - 1$. Notice in this expression how the finite differences and the stretching parameters conveniently group together. It is also apparent that because $\Delta \bar{w}_l$ is essentially the weighted average of the widths of cells l and $l - 1$, h_{wl} is the weighted average of e_{wl} and e_{wl-1} .

We can similarly expand the equations for the \hat{j} and \hat{k} terms which yields

$$\begin{aligned} & \left\{ \frac{\mu_p}{\mu_{x_{i,j+(1/2),k+(1/2)}}} \left[\frac{1}{(e_{yj} \Delta y_j)} (E z_{i,j+1,k+(1/2)}^s - E z_{i,j,k+(1/2)}^s) \right. \right. \\ & \left. \left. - \frac{1}{(e_{zk} \Delta z_k)} (E y_{i,j+(1/2),k+1}^s - E y_{i,j+(1/2),k}^s) \right] \right. \\ & \left. - \frac{\mu_p}{\mu_{x_{i,j+(1/2),k-(1/2)}}} \left[\frac{1}{(e_{yj} \Delta y_j)} (E z_{i,j+1,k-(1/2)}^s - E z_{i,j,k-(1/2)}^s) \right. \right. \\ & \left. \left. + \frac{1}{(e_{zk-1} \Delta z_{k-1})} (E y_{i,j+(1/2),k}^s - E y_{i,j+(1/2),k-1}^s) \right] \right\} \frac{1}{(h_{zk} \Delta z_k)} \\ & + \left\{ \frac{\mu_p}{\mu_{z_{i+(1/2),j+(1/2),k}}} \left[\frac{1}{(e_{yi} \Delta y_j)} (E x_{i+(1/2),j+1,k}^s - E x_{i+(1/2),j,k}^s) \right. \right. \\ & \left. \left. - \frac{1}{(e_{xi} \Delta x_i)} (E y_{i+1,j+(1/2),k}^s - E y_{i,j+(1/2),k}^s) \right] \right. \\ & \left. - \frac{\mu_p}{\mu_{z_{i-(1/2),j+(1/2),k}}} \left[\frac{1}{(e_{yj} \Delta y_j)} (E x_{i-(1/2),j+1,k}^s - E x_{i-(1/2),j,k}^s) \right. \right. \end{aligned}$$

$$\begin{aligned}
& - \frac{1}{(e_{xi-1} \Delta x_{i-1})} (Ey_{i,j+(1/2),k}^s - Ey_{i-1,j+(1/2),k}^s) \Bigg\} \\
& \cdot \frac{1}{(h_{xi} \Delta \bar{x}_i)} + i\omega \mu_p \hat{y}_p Ey_{i,j+(1/2),k}^s \\
& = -i\omega \mu_p (\hat{y}_{i,j+(1/2),k} - \hat{y}_p) Ey_{i,j+(1/2),k}^p \\
& - i\omega \mu_p \left\{ \left[\frac{(\mu_{x_{i,j+(1/2),k+(1/2)}} - \mu_p)}{\mu_{x_{i,j+(1/2),k+(1/2)}}} Hx_{i,j+(1/2),k+(1/2)}^p \right. \right. \\
& \quad \left. \left. - \frac{(\mu_{x_{i,j+(1/2),k-(1/2)}} - \mu_p)}{\mu_{x_{i,j+(1/2),k-(1/2)}}} Hx_{i,j+(1/2),k-(1/2)}^p \right] \frac{1}{(h_{zk} \Delta \bar{z}_k)} \right. \\
& \quad \left. - \left[\frac{(\mu_{z_{i+(1/2),j+(1/2),k}} - \mu_p)}{\mu_{z_{i+(1/2),j+(1/2),k}}} Hz_{i+(1/2),j+(1/2),k}^p \right. \right. \\
& \quad \left. \left. - \frac{(\mu_{z_{i-(1/2),j+(1/2),k}} - \mu_p)}{\mu_{z_{i-(1/2),j+(1/2),k}}} Hz_{i-(1/2),j+(1/2),k}^p \right] \frac{1}{(h_{xi} \Delta \bar{x}_i)} \right\} \quad (B6)
\end{aligned}$$

and

$$\begin{aligned}
& \left\{ \frac{\mu_p}{\mu_{y_{i+(1/2),j,k+(1/2)}}} \left[\frac{1}{(e_{zk} \Delta z_k)} (Ex_{i+(1/2),j,k+1}^s - Ex_{i+(1/2),j,k}^s) \right. \right. \\
& \quad \left. \left. - \frac{1}{(e_{xi} \Delta x_i)} (Ez_{i+1,j,k+(1/2)}^s - Ez_{i,j,k+(1/2)}^s) \right] \right. \\
& \quad \left. - \frac{\mu_p}{\mu_{y_{i-(1/2),j,k+(1/2)}}} \left[\frac{1}{(e_{zk} \Delta z_k)} (Ex_{i-(1/2),j,k+1}^s - Ex_{i-(1/2),j,k}^s) \right. \right. \\
& \quad \left. \left. - \frac{1}{(e_{xi-1} \Delta x_{i-1})} (Ez_{i,j,k+(1/2)}^s - Ez_{i-1,j,k+(1/2)}^s) \right] \right\} \frac{1}{(h_{xi} \Delta \bar{x}_i)} \\
& + \left\{ \frac{\mu_p}{\mu_{x_{i,j+(1/2),k+(1/2)}}} \left[\frac{1}{(e_{zk} \Delta z_k)} (Ey_{i,j+(1/2),k+1}^s - Ey_{i,j+(1/2),k}^s) \right. \right. \\
& \quad \left. \left. - \frac{1}{(e_{yj} \Delta y_j)} (Ez_{i,j+1,k+(1/2)}^s - Ez_{i,j,k+(1/2)}^s) \right] \right. \\
& \quad \left. - \frac{\mu_p}{\mu_{x_{i,j-(1/2),k+(1/2)}}} \left[\frac{1}{(e_{zk} \Delta z_k)} (Ey_{i,j-(1/2),k+1}^s - Ey_{i,j-(1/2),k}^s) \right. \right. \\
& \quad \left. \left. - \frac{1}{(e_{yi-1} \Delta y_{i-1})} (Ez_{i,j,k+(1/2)}^s - Ez_{i,j-1,k+(1/2)}^s) \right] \right\}
\end{aligned}$$

$$\begin{aligned}
& \cdot \frac{1}{(h_{yj} \Delta \bar{y}_j)} + i\omega \mu_p \hat{y}_p Ez_{i,j,k+(1/2)}^s \\
& = -i\omega \mu_p (\hat{y}_{i,j,k+(1/2)} - \hat{y}_p) Ez_{i,j,k+(1/2)}^p \\
& - i\omega \mu_p \left\{ \left[\frac{(\mu_{x_{i,j+(1/2),k+(1/2)}} - \mu_p)}{\mu_{x_{i,j+(1/2),k+(1/2)}}} Hx_{i,j+(1/2),k+(1/2)}^p \right. \right. \\
& \quad \left. \left. - \frac{(\mu_{x_{i,j-(1/2),k+(1/2)}} - \mu_p)}{\mu_{x_{i,j-(1/2),k+(1/2)}}} Hx_{i,j-(1/2),k+(1/2)}^p \right] \frac{1}{(h_{yj} \Delta \bar{y}_j)} \right. \\
& \quad \left. - \left[\frac{(\mu_{y_{i+(1/2),j,k+(1/2)}} - \mu_p)}{\mu_{y_{i+(1/2),j,k+(1/2)}}} Hy_{i+(1/2),j,k+(1/2)}^p \right. \right. \\
& \quad \left. \left. - \frac{(\mu_{y_{i-(1/2),j,k+(1/2)}} - \mu_p)}{\mu_{y_{i-(1/2),j,k+(1/2)}}} Hy_{i-(1/2),j,k+(1/2)}^p \right] \frac{1}{(h_{xi} \Delta \bar{x}_i)} \right\} \quad (B7)
\end{aligned}$$

respectively. Unfortunately the above equations will not produce a symmetrical matrix. Thus symmetrical scaling must be applied with (B5) being multiplied by $(e_{xi} \Delta x_i)(h_{yj} \Delta \bar{y}_j)(h_{zk} \Delta \bar{z}_k)$, (B6) by $(h_{xi} \Delta \bar{x}_i)(e_{yj} \Delta y_j)(h_{zk} \Delta \bar{z}_k)$, and (B7) by $(h_{xi} \Delta \bar{x}_i)(h_{yj} \Delta \bar{y}_j)(e_{zk} \Delta z_k)$.

Acknowledgments. We express our thanks to Ki Ha Lee of Lawrence Berkeley Laboratory for the use of the 1-D layered earth code used in the model comparisons. Discussions with Tarek Habashy of Schlumberger-Doll Research were helpful for successfully implementing variable magnetic permeability into the 3-D scheme. This work was performed at Sandia National Laboratories, which is operated for the U.S. Department of Energy. Funding for this project was provided by (1) DOE's office of Basic Energy Sciences, Division of Engineering and Geoscience, under contract DE-AC04-94AL85000; (2) the Sandia National Laboratory Director's Research funds; and (3) the VETEM project which is supported by the DOE Office of Technology and Development's "Landfill Stabilization Focus Area" program.

References

- Alumbaugh, D. L., and G. A. Newman, Fast, frequency domain electromagnetic modeling using finite differences, paper presented at 64th Annual International Meeting, Soc. Explor. Geophys., Los Angeles, Calif., Oct. 23-28, 1994.
- Alumbaugh, D. L., and G. A. Newman, Simulation of airborne electromagnetic measurements in three dimensional environments, paper presented at SAGEEP '95, Symposium on the Application of Geophysics to Environmental and Engineering Problems, Environ. and

- Eng. Geophys. Soc., Orlando, Fla., April 23–26, 1995.
- Barrett, R., M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. van der Vorst, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, Soc. for Ind. and Appl. Math., Philadelphia, 1994.
- Berenger, J., A perfectly matched layer for the absorption of electromagnetic waves, *J. Comput. Phys.*, **114**, 185–200, 1993.
- Chew, W. C., *Waves and Fields in Inhomogeneous Media*, Van Nostrand Reinhold, New York, 1994.
- Chew, W. C., and W. H. Weedon, A 3D perfectly matched medium from modified Maxwell's equations with stretched coordinates, *Microwave Opt. Tech. Lett.*, **7**, 599–604, 1994.
- Druskin, V., and L. Knizhnerman, A spectral semi-discrete method for the numerical solution of three-dimensional nonstationary problems of electric prospecting, *Izv. Earth Phys.*, **24**, 641–648, 1988.
- Druskin, V., and L. Knizhnerman, A spectral approach to solving three-dimensional diffusion Maxwell's equations in the time and frequency domains, *Radio Sci.*, **29**, 937–953, 1994.
- Freund, R., Conjugate gradient type methods for linear systems with complex symmetric coefficient matrices, *SIAM J. Sci. Stat. Comput.*, **13**, 425–448, 1992.
- Habashy, T. M., R. W. Groom, and B. R. Spies, Beyond the Born and Rytov approximations: A nonlinear approach to electromagnetic scattering, *J. Geophys. Res.*, **98**, 1759–1775, 1993.
- Katz, D. S., E. T. Thiele, and A. Taflove, Validation and extension to three dimensions of the Berenger PML absorbing boundary condition for FD-TD meshes; *IEEE Micro. Guided Wave Lett.*, **4**, 268–270, 1994.
- Newman, G. A., Cross well electromagnetic inversion using integral and differential equations, *Geophysics*, **60**, 899–911, 1995.
- Newman, G. A., and D. L. Alumbaugh, Frequency-domain modeling of airborne electromagnetic responses using staggered finite differences, *Geophys. Prospect.*, **43**, 1021–1042, 1995.
- Newman, G. A., G. W. Hohmann, and W. L. Anderson, Transient electromagnetic response of a three-dimensional body in a layered earth, *Geophysics*, **51**, 1608–1627, 1986.
- Pellerin, L., V. F. Labson, and M. C. Pfeifer, VE-TEM—a very early time electromagnetic system, paper presented at SAGEEP '95, Symposium on the Application of Geophysics to Environmental and Engineering Problems, Environ. and Eng. Geophys. Soc., Orlando, Fla., April 23–26, 1995.
- Sarkar, T. K., On the application of the generalized biconjugate gradient method, *J. Electromagn. Waves Appl.*, **1**, 223–242, 1987.
- Skjellum, A., N. E. Doss, and P. V. Bangalore, Writing libraries in MPI, in *Proceedings of the Scalable Parallel Libraries Conference*, edited by A. Skjellum and D. Reese, pp. 166–173, 1993. IEEE Comput. Sci. Press.
- Smith, T. J., Conservative modeling of 3-D electromagnetic fields, paper presented at 11th Workshop on Electromagnetic Induction in the Earth, Int. Assoc. of Geomagn. and Aeron., Wellington, New Zealand, Aug. 26–Sept. 2, 1992.
- Stewart, D. C., W. L. Anderson, T. P. Grover, and V. F. Labson, Shallow subsurface mapping by electromagnetic sounding in the 300 kHz to 30 MHz range: Model studies and prototype system assessment, *Geophysics*, **59**, 1201–1210, 1994.
- Torres-Verdín, C., and T. M. Habashy, Rapid 2.5-D forward modeling and inversion via a new nonlinear scattering approximation, *Radio Sci.*, **29**, 1051–1079, 1994.
- Tripp, A. C., and G. W. Hohmann, Block diagonalization of the electromagnetic impedance matrix of a symmetric buried body using group theory, *IEEE Trans. Geosci. Remote Sens.*, **GE-22**, 62–69, 1984.
- Wilt, M., H. F. Morrison, A. Becker, H. W. Tseng, K. H. Lee, C. Torres-Verdin, and D. Alumbaugh, Crosshole electromagnetic tomography: A new technology for oil field characterization, *Leading Edge*, **14**, 173–177, 1995.
- Wang, T., and G. W. Hohmann, A finite difference time-domain solution for three-dimensional electromagnetic modeling, *Geophysics*, **58**, 797–809, 1993.
- Xiong, Z., Electromagnetic modeling of 3-D structures by the method of system iteration using integral equations, *Geophysics*, **57**, 1556–1561, 1992.
- Yee, K. S., Numerical solution of initial boundary problems involving Maxwell's equations in isotropic media, *IEEE Trans. Antennas Propag.*, **AP-14**, 302–309, 1966.
- Zhdanov, M. S., and S. Fang, Quasi linear approximation in 3-D electromagnetic modeling, paper presented at the International Symposium on Three-Dimensional Electromagnetics, Schlumberger-Doll Res., Ridgefield, Conn., Oct. 4–6, 1995.
- D. L. Alumbaugh, G. A. Newman, L. Prevost, and J. N. Shadid, Sandia National Laboratory, P. O. Box 5800, M.S. 0750, Albuquerque, NM 87185.

(Received April 12, 1995; revised August 17, 1995; accepted September 11, 1995.)